arXiv:1603.03815v3 [q-bio.QM] 24 Jan 2017

Initial validation for the estimation of resting-state fMRI effective connectivity by a generalization of the correlation approach☆

Nan Xu[a,*], R. Nathan Spreng[b], Peter C. Doerschuk[a,c]

[a]*School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, USA*
[b]*Human Neuroscience Institute, Cornell University, Ithaca, NY, USA*
[c]*Nancy E. and Peter C. Meinig School of Biomedical Engineering, Cornell University, Ithaca, NY, USA*

Abstract

Resting-state functional MRI (rs-fMRI) is widely used to noninvasively study human brain networks. Network functional connectivity is often estimated by calculating the timeseries correlation between blood-oxygen-level dependent (BOLD) signal from different regions of interest. However, standard correlation cannot characterize the direction of information flow between regions. In this paper, we introduce and test a new concept, prediction correlation, to estimate effective connectivity in functional brain networks from rs-fMRI. In this approach, the correlation between two BOLD signals is replaced by a correlation between one BOLD signal and a prediction of this signal via a causal system driven by another BOLD signal. Three validations are described: (1) Prediction correlation performed well on simulated data where the ground truth was known, and outperformed four other methods. (2) On simulated data designed to display the "common driver" problem, prediction correlation did not introduce false connections between non-interacting driven ROIs. (3) On experimental data, prediction correlation recovered the previously identified network organization of human brain. Prediction correlation scales well to work with hundreds of ROIs, enabling it to assess whole brain interregional connectivity at the single subject level. These results provide an initial validation that prediction correlation can capture the direction of information flow and estimate the duration of extended temporal delays in information flow between regions of interest based on BOLD signal. This approach not only maintains the high sensitivity to network connectivity provided by the correlation analysis, but also

☆The authors certified that they do have conflict of interest.

*Corresponding author
Email addresses: nx25@cornell.edu (Nan Xu), nathan.spreng@gmail.com (R. Nathan Spreng), pd83@cornell.edu (Peter C. Doerschuk)

performs well in the estimation of causal information flow in the brain.

Keywords: effective connectivity, functional connectivity, functional connected network, resting-state fMRI, correlation analysis

---

1. Introduction

Resting-state functional MRI (rs-fMRI) has been widely used to study the intrinsic functional architecture of the human brain based on spontaneous oscillations of the blood oxygen level dependent (BOLD) signals (Smith et al., 2011; Biswal et al., 1995; Yeo et al., 2011; Power et al., 2011). One fruitful approach has been to examine the correlations between rs-fMRI timeseries at pairs of regions of interest (ROIs) and use the correlations as a measure of connectivity strength between each pair (Wig et al., 2011; Sporns, 2011). The correlation method, though simple, plays a fundamental role in evaluating functional connectivity in the human brain for both task-evoked networks (Sadaghiani et al., 2015; Cole et al., 2014) and resting-state networks (Hipp and Siegel, 2015; Sadaghiani et al., 2015; Power et al., 2013). The relationships between correlation and the topological properties, including small-world organization, modular structure, and highly connected hubs, has been studied in Zalesky et al. (2012). However, the direction of information flow between pairs of ROIs and the causality of information flow cannot be derived from standard correlation methods. Reliable insight into the direction and causality of functional connections in the brain from BOLD signals would provide substantial breakthroughs in characterizing large-scale brain network dynamics.

The BOLD signal is an indirect and sluggish measure of neuronal activity. Despite this, substantial insights have been gleaned by examining patterns of BOLD signals as proxies for functional connectivity in the brain, and these are consistent with more direct and invasive observations (Foster et al., 2015). At every level of analysis, the brain demonstrates an organized network structure (Bassett and Gazzaniga, 2011). So, even though neuronal activation occurs on the millisecond time scale, organized and structured activation patterns are also observed on the level of seconds, which is within the range of BOLD signals and is important for understanding cognition. Causal information about the flow of information in the brain may be detected and estimated from the BOLD signals. It remains critical, however, to evaluate methods of investigation against ground truth simulation in order to validate these methods.

Numerous methods for estimating functional or effective connectivity (Van Den Heuvel and Pol, 2010; Friston, 2011) have recently been evaluated against ground truth networks using simulated rs-fMRI data (Smith et al., 2011). Functional connectivity can be quantified with a measure of statistical dependence such as correlation, whereas effective connectivity measures the directed causal influence (Friston, 2011). In Smith et al. (2011), performance of both types of methods across a range of measures was mixed. Standard and partial correlation excelled at detecting the presence of a connection. Other methods for estimating the direction of a connection varied from chance (Granger) to greater than 50% accuracy (Patel's Tau and pairwise LiNGAM). These results suggest that novel methods are needed to estimate directed connectivity from rs-fMRI data, particularly with a large number of ROIs, which are necessary for full coverage of cortical and subcortical areas in the human brain. In this paper, we introduce a new method, prediction correlation, to the neuroimaging community and provide an initial validation of the approach.

Methods for estimating functional connectivity can be oriented toward estimating a real number describing strength of connectivity, which might be quite small, versus estimating a binary connectivity, which is present or absent, with possibly the addition of a strength of connectivity, in the form of a real number, for the case where a connection is present. Correlation and prediction correlation, which is a generalization of correlation that we propose in this paper, are methods that estimate a real number that describes strength of connection. Subsequent processing can then be applied to remove weak connections and/or organize the complete network into modular networks.

As is described in the following sections, testing on simulated rs-fMRI data with known ground-truth networks (Smith et al., 2011) demonstrates that prediction correlation is not only sensitive in detecting network connections, as identified by standard correlation, but also achieves the highest accuracy on estimation of connection directionality among all approaches used in Smith et al. (2011) (Section 3.1). In a "common driver" phenomena, when ROI 1 drives ROIs 2 and 3 but ROIs 2 and 3 do not directly interact, prediction correlation correctly detects strong 1→2 and 1→3 connections but not 2→3 or 3→2 connections (Section 3.2). Finally, extending Xu et al. (2014), we demonstrate the robustness of this method on experimental data and that prediction correlation recovers previously identified brain network organization from experimental data (Section 3.3).

2. Methods: Prediction correlation

2.1. Fundamental method

In what follows, we describe a methodology for analyzing rs-fMRI data using a generalization of the well-established correlation approach, which is to correlate the timeseries at two ROIs. The generalization, denoted by "p-correlation" ("p" for "prediction") is to replace correlation between the BOLD timeseries at two ROIs by correlation between the BOLD timeseries at one ROI and a prediction of this timeseries. The prediction is the output of a mathematical dynamical system that is driven by the timeseries at the other ROI. More generally, the prediction could be based on several, spatially discrete, ROIs. In this paper, we focus on the case where only one other ROI is used. We assume that the dynamical system is linear and has finite memory and that the memory duration and parameters may be estimated from the BOLD timeseries. If the prediction of the timeseries is restricted to use only the current value of the timeseries that drives the dynamical system, then p-correlation is the same as standard correlation. Therefore p-correlation is a generalization of correlation. Features of p-correlation include (1) the ability to indicate the directionality of the interaction between two ROIs, (due to the fact that this prediction correlation is asymmetrical between two signals), and (2) the ability to evaluate the interaction based on casual information.

In the remainder of this section, we describe the p-correlation approach in detail. Consider the ordered pair of ROIs (i, j) and let x_i (x_j) denote the rs-fMRI timeseries at the i^{th} (j^{th}) ROI. Both timeseries have duration N_x. The x_j signal is predicted from the x_i signal by a linear time-invariant causal dynamical model with x_i as the input and the prediction $\hat{x}_{j|i}$ as the output. This model can be described by an impulse response, denoted $h_{j|i}$, which is zero for negative times. We assume that the impulse response is of finite duration, with duration denoted by $N_{h_{j|i}}$. In summary,

$$\hat{x}_{j|i}[n] = \sum_{m=0}^{N_{h_{j|i}}} h_{j|i}[m] x_i[n-m]. \tag{1}$$

The basic approach to estimate the coefficients of $h_{j|i}$ is to minimize the least squares cost

$$\mathcal{J}(h_{j|i}) = \sum_{n=0}^{N_x-1} (x_j[n] - \hat{x}_{j|i}[n])^2. \tag{2}$$

We estimate the value of $N_{h_{j|i}}$ and the values of the impulse response at the same time by restating the least squares problem as a Gaussian maximum likelihood estimator (MLE) with a known variance for the measurement errors. The MLE allows a trade off of the accuracy of predicting the current data (i.e., minimizing $\mathcal{J}$), which is best done by large values of $N_{h_{j|i}}$, with the accuracy of predicting when presented with new data, which is best done by smaller values of $N_{h_{j|i}}$. There are several approaches to quantifying this trade off including Akaike information criteria (AIC) (Akaike, 1974, 1970; Sugiura, 1978; Hurvich and Tsai, 1989, 1993; Cavanaugh, 1997), Bayesian information criteria (BIC) (Schwarz, 1978), restricted maximum likelihood (REML) (Thompson Jr, 1962; Patterson and Thompson, 1971), minimum description length (Rissanen, 1978) and minimum message length (Wallace and Boulton, 1968). We have focused on AIC because it leads to easily computed problem formulations (Eq. 3). AIC realizes this balancing goal by minimizing the sum of two terms, one term that characterizes the prediction error of the dynamic system through the least squares cost $\mathcal{J}(h_{j|i})$ and a second term that depends on the durations $N_{h_{j|i}}$ and N_x:

$$\text{AIC} = \begin{cases} N_x \log(\frac{2\pi}{N_x - N_{h_{j|i}}} \mathcal{J}(h_{j|i})) + N_x + N_{h_{j|i}} & \text{if } N_x/N_{h_{j|i}} \geq 40 \\ N_x \log(\frac{2\pi}{N_x - N_{h_{j|i}}} \mathcal{J}(h_{j|i})) + N_x^2 + \frac{N_{h_{j|i}}^2 - N_x + N_{h_{j|i}}}{N_x - N_{h_{j|i}} - 1} & \text{otherwise} \end{cases} . \quad (3)$$

Simultaneous minimization of Eq. 3 with respect to both $h_{j|i}$, which occurs only in the $\mathcal{J}(h_{j|i})$ term, and $N_{h_{j|i}}$ determines the duration and the value of the impulse response. The integer minimization over $N_{h_{j|i}}$ is computed by testing each value in a predetermined range of values, i.e., 1,2, ..., D seconds. Then, for each value of $N_{h_{j|i}}$, the minimization with respect to $h_{j|i}$ involves only minimizing $\mathcal{J}(h_{j|i})$. Since the dynamical system describing how x_i influences x_j is separate from the dynamical system describing how x_j influences x_i, the approach described here can lead to a directed rather than undirected graph of interactions between ROIs.

Once $h_{j|i}$ and $N_{j|i}$ are estimated, the output of the dynamical system, which is the prediction $\hat{x}_{j|i}$, can be computed, and then the correlation of x_j and $\hat{x}_{j|i}$, which is the so-called p-correlation, denoted by $\rho_{j|i}$, can be computed. We use "correlation" and $\rho_{j,i}$ for the standard approach (i.e., the standard correlation between x_j and x_i).

Let the total number of ROIs be denoted by N_{ROI}. P-correlation is an asymmetric $N_{\text{ROI}} \times N_{\text{ROI}}$ matrix, where the asymmetry follows from $\rho_{j|i} \neq \rho_{i|j}$. Furthermore, p-correlation includes lags of the x_i signal since the dynamical system output at time n,

$\hat{x}_{j|i}[n]$, depends on the input at its current and previous times, i.e., $x_i[n]$, $x_i[n-1]$, ..., $x_i[n - N_{h_{j|i}} + 1]$. If $N_{h_{j|i}} = 1$ (i.e., no lags) and $h_{j|i}[0] \geq 0$ then $\rho_{j|i}$ is the correlation between x_j and x_i so that $\rho_{j|i} = \rho_{j,i}$ and the approach of this paper exactly reduces to the standard approach. In Section 2.2.1, we describe a constraint such that $h_{j|i}[0] \geq 0$ is always achieved. The p-correlation method does not depend upon the sampling rate (TR) which allows for collapsing across different scan sites or studies. The entire algorithm is shown in Fig. 1. Matlab software implementing p-correlation is available upon request.

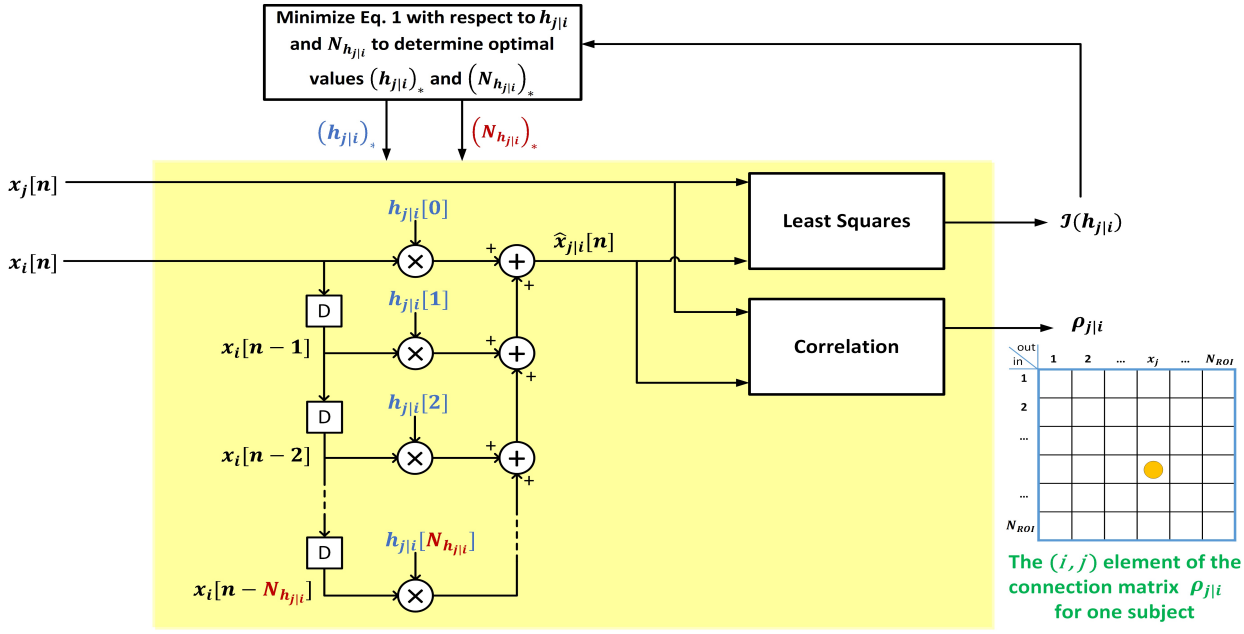

(a) Flow chart describing the computational method.

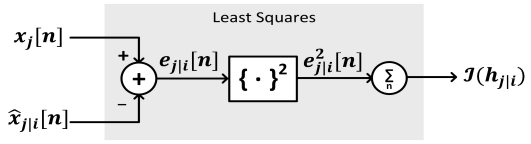

(b) "Least Squares" Sub-block.

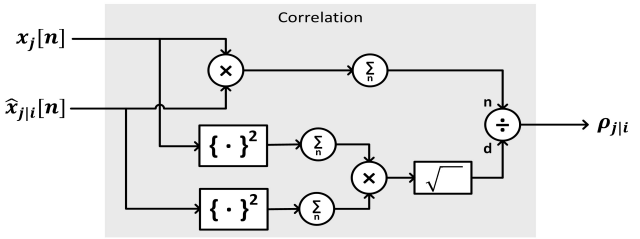

(c) "Correlation" Sub-block.

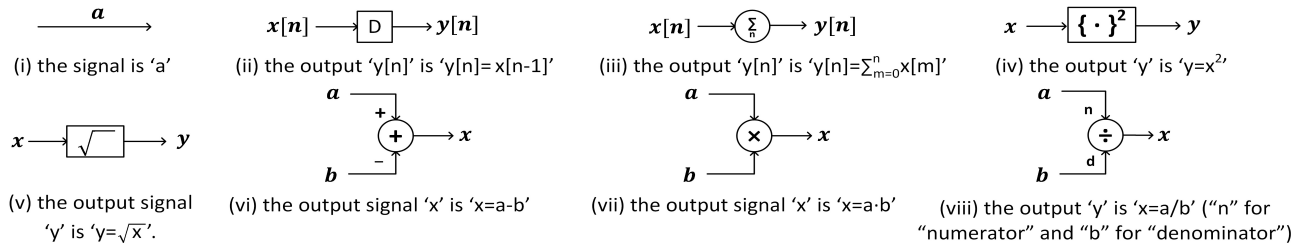
a
(i) the signal is 'a'
x[n] D y[n]
(ii) the output 'y[n]' is 'y[n]= x[n-1]'
x[n] $\sum_n$ y[n]
(iii) the output 'y[n]' is 'y[n]=$\sum_{m=0}^{n}$x[m]'
x $\{\cdot\}^2$ y
(iv) the output 'y' is 'y=x^2'
x y
(v) the output signal 'y' is 'y=$\sqrt{x}$'.
a b x
(vi) the output signal 'x' is 'x=a-b'
a b x
(vii) the output signal 'x' is 'x=a·b'
a b n d x
(viii) the output 'y' is 'x=a/b' ("n" for "numerator" and "b" for "denominator")

(d) Operators used in the block diagrams of Panels (a)-(c). Lines that intersect in the form —⌒—, do not connect at the intersection while other intersections imply connections.

Figure 1: Block diagram and sub-block diagrams describing the computation of p-correlation for one pair of ROIs.

2.2. Specializations of the fundamental method

In Section 2.1 we defined p-correlation and described a practical method for its computation. The result is an asymmetric matrix of connection strengths for each subject. This fundamental method can be specialized for particular applications, often based on user's interests and what the user knows about the details of the applications. Several such specializations are described in the following paragraphs.

2.2.1. Constraints on the least squares problems

If the user has information on the type of interactions that are present, then this information can be used as a constraint on the least squares problem that determines the impulse response which is the basis of the prediction. For example, as in the simulated data of Smith et al. (Smith et al., 2011), the interactions are all positive. Constraining the impulse response values $h_{j|i}[n]$ to be nonnegative has implications for the values of $\rho_{j|i}$. Let $R_{j|i}$ be the covariance of x_j and $\hat{x}_{j|i}$. $R_{j|i}$ is related to the covariance of $x_j[n]$ and $x_i[n-m]$ (i.e., the m-lagged covariance of the two signals, denoted by $R_{j,i}[m]$) by $R_{j|i} = \sum_{m=0}^{N_{h_{j|i}}-1} R_{j,i}[m]h_{j|i}[m]$. The covariance $R_{j|i}$ is the numerator of $\rho_{j|i}$. Therefore, if all the lagged covariances are positive and we require the estimated values of $h_{j|i}[m]$ to be positive then we are assured of getting a nonnegative value for $R_{j|i}$ and for the p-correlation $\rho_{j|i}$. In the traditional functional connectivity analysis, when global signal regression is applied to rs-fMRI timeseries data, the valid inference of negative correlations cannot be made (Murphy et al., 2009; Saad et al., 2012), and only positive correlations are interpreted. In this situation, the nonnegative "constrained" estimation approach is appropriate.

2.2.2. Thresholding $\rho_{j|i}$

Three natural methods for thresholding $\rho_{j|i}$ are described in this section.

Even with $h_{j|i}[n] \geq 0$, it may be that p-correlation is not positive because one or more of the m-lagged covariance values are negative. Therefore, if non-negativity is required, we replace all negative $\rho_{j|i}$ values by zeros. One reason for seeking to have $\rho_{j|i}$ non negative is mean signal regression in the preprocessing of the fMRI data which makes it difficult to interpret negative correlations. However, alternative preprocessing which omits mean signal regression (Jo et al., 2013) removes this requirement.

The previous paragraph concerned thresholding at value 0. Higher data-dependent minimum thresholds are often used for correlation and the same approach can be applied to

p-correlaton. A standard approach (Power et al., 2011) is to order the values of correlation and leave the top s percent of values unchanged and set the remaining values to zero. In other words, the threshold $\gamma(s)$ is set to be the 100-s percentile of all values in the p-correlation matrix.

In some problems the interactions are known to be unidirectional, e.g., in the simulated data of Smith (Smith et al., 2011). In this situation, a third thresholding method, which makes p-correlation unidirectional, is natural. The threshold is to consider the two transpose-related elements of the matrix and set the smaller to zero and leave the larger unchanged.

Each of the thresholding methods is a nonlinear operation applied to the matrix of $\rho_{j|i}$ coefficients. Each can be applied to any matrix M to give an output matrix N, in particular, in the order of the previous three paragraphs,

$$N_{ij} = \begin{cases} M_{ij}, & \text{if } M_{ij} \geq 0 \\ 0, & \text{otherwise} \end{cases}, \tag{4a}$$

$$N_{ij} = \begin{cases} M_{ij}, & \text{if } M_{ij} \geq \gamma(s) \\ 0, & \text{otherwise} \end{cases}, \tag{4b}$$

where $\gamma(s)$ is the $100 - s$ percentile of all values in M, and

$$N_{ij} = \begin{cases} M_{ij}, & \text{if } M_{ij} \geq M_{ji} \\ 0, & \text{otherwise} \end{cases}. \tag{4c}$$

The thresholding approach forms a $N_{\text{ROI}} \times N_{\text{ROI}}$ matrix of thresholded connection weights, from which the network is computed.

2.2.3. Averaging over subjects

Some investigations, e.g., Smith et al. (2011); Laumann et al. (2015), are interested in estimating subject-by-subject details, but in many other investigations on functional networks of human brain using experimental data, e.g., Power et al. (2011, 2013); Schaefer et al. (2014); Gordon et al. (2016), there is averaging over subjects in order to improve the SNR. Just as the thresholding methods (Section 2.2.2), which are nonlinearities that can be applied to any matrix, the averaging we use can be applied to any family of matrices M_k ($k \in \{1, \ldots, K\}$, where K is the number of subjects) to give an output matrix N via

$N = \frac{1}{K}\sum_{k=1}^{K} M_k$. The functional network estimated by the averaged p-correlation matrix can be further clustered into sub-networks through a graphic theoretic analysis.

2.3. Extension to multi-subject processing

There is a recent interest in estimating effective networks from multiple subjects while accommodating the heterogeneity of the group (Smith, 2012; Ramsey et al., 2010; Gates and Molenaar, 2012). Specifically, the IMaGES algorithm (Ramsey et al., 2010) estimates one generalized network from a group by assuming all subjects are homogeneous, and the GIMME algorithm (Gates and Molenaar, 2012) can further refine the estimate for each individual subject from the general information estimated from the whole group. IMaGES and GIMME are based on existing single-subject methods, specifically GES for IMaGES and uSEM and euSEM for GIMME and, when applied to groups of appropriate size, both GIMME and IMaGES provide more accurate estimates of effective connectivity than the single subject methods on which they are based (Ramsey et al., 2011; Gates and Molenaar, 2012).

Information concerning groups of subjects could also be used in p-correlation. One approach would be to replace the $h_{j|i}$ in Eq. 1 by $h^g_{j|i} + h^l_{j|i}$, where $h^g_{j|i}$ is the group component common to all subjects, and $h^l_{j|i}$ is the component unique to the specific subject l. In this approach, Eq. 1 would be generalized to

$$\hat{x}^l_{j|i}[n] = \sum_{m=0}^{N_{h^g_{j|i}}} h^g_{j|i}[m] x^g_i[n-m] + \sum_{k=0}^{N_{h^l_{j|i}}} h^l_{j|i}[k] x^l_i[n-k]^l \tag{5}$$

where $N_{h^g_{j|i}}$ and $N_{h^l_{j|i}}$ are the probably different durations of the two components of the causal finite-duration impulse response. There are two issues when using Eq. 5. First the AIC analysis must be generalized in order to determine two impulse response durations where one is common to the entire group of subjects. Second, in order to require the least squares to use the group impulse response and not just set it to zero, a regularizer such as $\sum_{m=0}^{N_{h^l_{j|i}}} (h^l_{j|i}[m])^2$ must be added to the least squares cost. While both of these issues can be addressed, in the current paper, we only focus on the individual analysis, which may be the only meaningful option under certain circumstances, i.e., a clinical environment.

3. Results

3.1. Application on Simulated Data

3.1.1. Data source: simulated BOLD timeseries

Simulated fMRI timeseries from the laboratory of S. M. Smith are documented (Smith et al., 2011) and available on-line (`http://www.fmrib.ox.ac.uk/analysis/netsim/`). These timeseries have been used as benchmark simulated fMRI data for testing effective connectivity (Smith et al., 2011; Gates and Molenaar, 2012; Ramsey et al., 2011; Hyvärinen and Smith, 2013). The simulations are based on a variety of underlying networks of different complexity and can be described as having three levels. First there is a neural level which is a stochastic linear vector differential equation which produces a neural timeseries for each ROI. Second, for each ROI, there is a nonlinear balloon model driven by the corresponding neural timeseries which produces a vascular timeseries. Third, for each ROI, the fMRI timeseries is the vascular timeseries plus thermal noise. To simulate preprocessing of fMRI data, a high-pass filtered at a cutoff frequency of 1/200s was applied to each simulation (most recently revised on Aug. 24, 2012 based on the website `www.fmrib.ox.ac.uk/analysis/netsim`). The current paper considers the first four sets of simulations from Smith et al. (2011), $Sim1 - Sim4$, which are the four most "typical" network scenarios provided in Smith et al. (2011), and which are based on different underlying networks with sizes 5, 10, 15, and 50 ROIs, respectively.

These synthetic fMRI timeseries were sampled every 3s (TR= $3s$) and the total duration is $N_x = 10$ mins. All four simulations have 1% thermal noise and the hemodynamic response function (HRF) used in the second step has standard deviation of 0.5 s. The simulation is repeated for each of 50 subjects.

3.1.2. Specialization on p-correlation for the processing of the simulated data

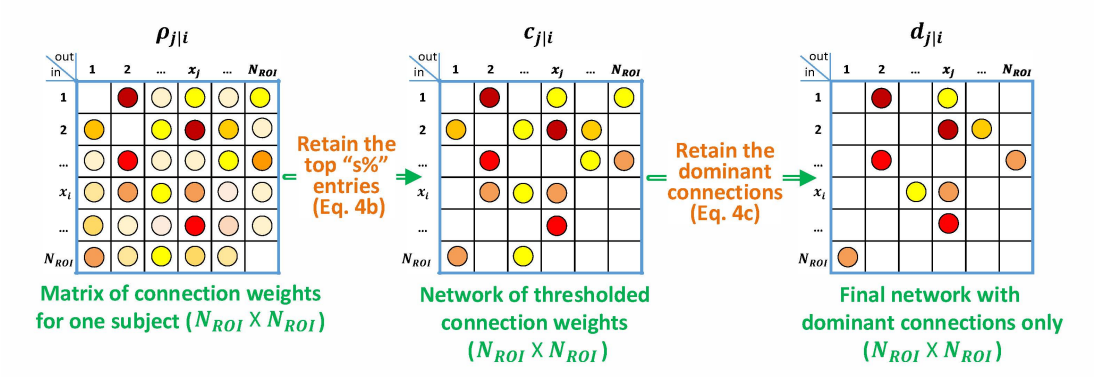

Figure 2: Block diagram describing the specialization of p-correlation for simulated data. Nonzero entries are filled by colored dots with higher values represented by "hotter" colors and lower values represented by "colder" colors, and zero entries are left as blank in the above matrices.

The algorithm is shown in Figure 2. Given that the interactions are all positive in the simulated data, it is natural to apply the nonnegative constraint on the least squares problem so that no negative impulse responses are allowed. Although unconstrained p-correlation is also computed on the simulated data, looking forward to Section 3.1.5, the numerical results indicate that the constrained version is more appropriate.

As is described above, the integer minimization over the impulse function duration, $N_{h_{j|i}}$, is computed by testing from 1 second up to D seconds. Assuming that knowledge of the behavior of a ROI over the past 15 seconds is sufficient to describe its effect on a second ROI, we restricted the temporal window for directional influence between ROIs to no more than 15s, i.e. $D = 15s$.

Next, we consider the choice of threshold, s in Eq. 4b. We use this method in order to exploit all of the a priori knowledge about the simulated data. Since the underlying ground truth networks for the simulated fMRI timeseries, denoted by $a_{j|i}$, are given, the threshold value s is among our prior knowledge as is described below. We denote ROIs that are involved in the connections of the ground truth network as active ROIs. All connections involving the active ROIs are connections of interest (COIs), including connections that are actually absent such as the reverse connection in an unidirectional interaction. The value of s is then the ratio of the number of COIs and the number of all possible connections, which

gives $s = 40,\ 22,\ 16$ and 4 percent for the four simulations, respectively. An example of computing s for a 5-node network is shown in Fig. 3.

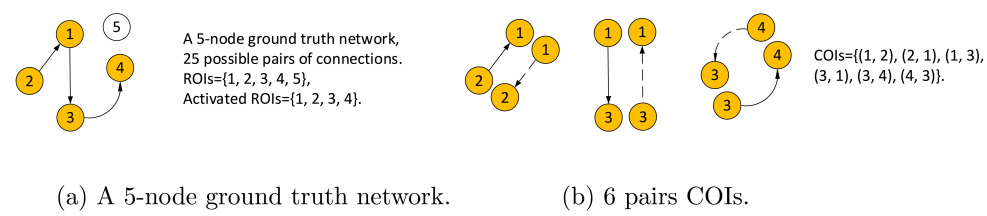

(a) A 5-node ground truth network. (b) 6 pairs COIs.

Figure 3: Example calculation of the threshold s for a 5-node network. (a) The network with activated ROIs shown in orange. The number of all possible connections is $5^2 = 25$. (b) The 6 COIs, where the dashed lines are connections that do not existed in the ground truth but still are considered interesting. Therefore, $s = 6/25$ for this network.

For the Smith simulated data, we have additional prior knowledge that the networks contain only unidirectional connections. Therefore, as is also done in Smith et al. (2011), we compare our estimated network $d_{j|i}$, which includes the unidirectional condition, with the ground truth network $a_{j|i}$. The estimated network $d_{j|i}$ is the output of Eq. 4c where the input is the thresholded network $c_{j|i}$.

3.1.3. Performance Criteria

To compare the computed and ground truth networks, we define "accuracy", denoted by $\mathcal{A}$. In particular, $\mathcal{A}$ is defined to be the mean fractional rate of detecting the correct directionality of true connections. Specifically, it is defined to be

$$\mathcal{A} = \frac{\sum_{i=1}^{N_{\mathrm{ROI}}} \sum_{j=1}^{N_{\mathrm{ROI}}} 1\{a_{j|i} > 0\} 1\{d_{j|i} > 0\}}{\sum_{i=1}^{N_{\mathrm{ROI}}} \sum_{j=1}^{N_{\mathrm{ROI}}} 1\{a_{j|i} > 0\}}, \tag{6}$$

where $1\{L\}$ is 1 if L is true, and 0 otherwise. Like the computation of the "d-accuracy" introduced in Smith et al. (2011), $\mathcal{A}$ evaluates the percentage of the correct directionality ($\mathcal{A}$ is between 0 and 1). The threshold operation introduced above (Section 3.1.2) differentiates the performance of directional analytical methods based on their sensitivity. The more sensitive the method is, the more true connections it can detect. Notice that application of the threshold s leads to $d_{j|i}$ values that are almost certainly far from zero or exactly zero. Computing the accuracy $\mathcal{A}$ after the threshold operation tells the directionality after know-

ing the presence of the connections, which enables us to evaluate the overall performance of sensitivity and directionality of a directional analytical method.

3.1.4. Alternative methods for effective networks estimation

P-correlation and four alternative methods from Smith et al. (2011), specifically, "Granger B1", "Gen Synch S1", "LiNGAM" and "Patel's conditional dependence measure", were compared by the accuracy criteria ($\mathcal{A}$), since under both synthetic and experimental scenarios, these methods have been tested and have relatively good performances among all the others (Smith et al., 2011; Dawson et al., 2013). The computation of these methods were done by software provided by Prof. S.M. Smith. Granger B1, a pairwise Granger causality estimation method which provides the best performance among Granger causality approaches (Smith et al., 2011; Dawson et al., 2013), uses the Bayesian Information Criterion to estimate the lag up to 1 TR. Gen Synch S1 is a nonlinear synchronization method with respect to the time lag 1 TR. It "evaluates synchrony by analyzing the interdependence between the signals in a state space reconstructed domain" (Dauwels et al., 2010, p. 671). The LiNGAM (Linear, Non-Gaussian, Acyclic causal Models) algorithm is a global network model utilizing higher-order distributional statistics, via independent component analysis, to estimate the network connections. Patel's conditional dependence measure investigates the causality from the imbalance between two conditional probabilities, $P(x_j|x_i)$ and $P(x_i|x_j)$. P-correlation, Granger B1, Gen Synch S1 and LiNGAM all compute an asymmetric matrix filled with real-number connection weights, analogous to our $c_{j|i}$. In all cases, the unidirectional prior knowledge is applied analogous to our transformation from $c_{j|i}$ to $d_{j|i}$. For the Patel method implemented by Smith et al. (2011), the thresholding operation was applied on "Patel's κ bin 0.75" matrix, while the directionality was determined by "Patel's τ bin 0.75" matrix.

In addition to the algorithms included in Smith et al. (2011), IMaGES (Ramsey et al., 2010) and uSEM (Kim et al., 2007) which is the estimation method for resting-state fMRI employed by GIMME algorithm, have also been tested on the same set of simulated data (Ramsey et al., 2011; Gates and Molenaar, 2012). Results reported in Ramsey et al. (2011); Gates and Molenaar (2012) show that their estimation based on the single subject is either similar to or less good than the best-performing method provided in Smith et al. (2011).

Comparing p-correlation with alternative methods of estimating effective connectivity, p-correlation provides a full asymmetric matrix for each subject independent of all other sub-

jects, in which each entry, like correlation, predicts a connection strength between two ROIs. The ability to compute results based on an individual subject means that p-correlation can potentially be used in a clinical environment. This full asymmetric matrix of p-correlations can be thresholded as desired and/or further processed as desired using another algorithm, i.e., a graph analytic algorithm. In addition, p-correlation can process networks with hundreds of ROIs while GIMME is limited to 3-25 ROIs (Page 3 of GIMME Manual (Version 12)). Furthermore, p-correlation estimates the temporal causal relation in the form of lagged impulse response in addition to the spatial causal relation between any pair of ROIs. In contrast, some alternative algorithms (e.g. IMaGES) estimate a sparse graph of interactions, and thus solve a somewhat different problem than the p-correlation method. Other algorithms have been developed as post-processing algorithms, which cannot detect connections, but only estimate direction if connections are detected by other methods, e.g., correlation. Among them, pairwise LiNGAM (Hyvärinen and Smith, 2013) achieved success on Smith's data (Smith et al., 2011). Several algorithms, such as Patel's τ, LiNGAM and pairwise LiNGAM, chose one of the two possible directions for each pair of ROIs. Such unidirectionality may be appropriate in some situations. Alternative algorithms, including p-correlation, provide strengths for both directions, where the two strengths may be quite different when one direction is dominant.

3.1.5. Results on simulated data

The methods described in this paper were implemented in Matlab software, which is available upon request, and were applied to four of Smith's fMRI simulations (Smith et al., 2011). The four simulations are $Sim1 - Sim4$ which have a variable number of ROIs (5, 10, 15, 50) but no confounding variables.

The p-correlation method is based on estimation of a linear time-invariant causal dynamic model. The sample means of the duration of either constrained or unconstrained impulse responses are 3.34s, 3.58s, 3.64s and 3.76s for the 4 simulations, respectively. By limiting the impulse response duration to 1 TR, it was verified that p-correlation with constraint on Least Squares is equivalent to the standard correlation as is described in Section 1. After thresholding the p-correlations computed with the nonnegative constraint on the coefficients of the linear system, an asymmetric matrix of connection weights $c_{j|i}$ for each subject was obtained.

The same specifications for processing of the simulated data, in particular, the same

choice of the s threshold (Eq. 4b) and the knowledge of unidirectionality (Eq. 4c), have also been applied to the results of four alternative methods introduced in Section 3.1.4. The performance of all five methods was evaluated by the accuracy criteria $\mathcal{A}$ (Eq. 6) for each subject. Fig. 4 shows the input to the accuracy criteria $\mathcal{A}$, i.e., $a_{j|i}$ and $d_{j|i}$, for Subject 14 of $Sim2$.

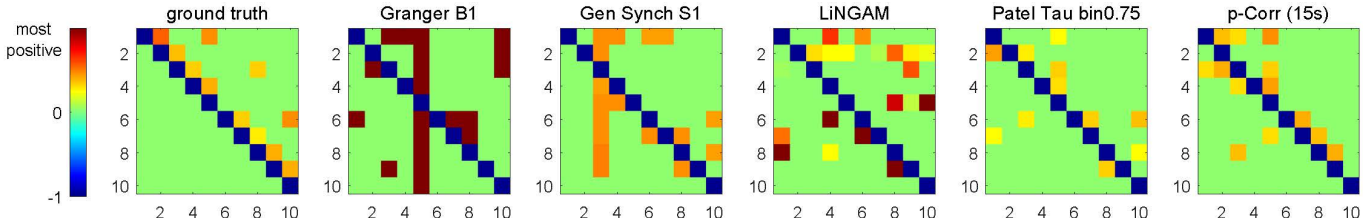

Figure 4: Images of $a_{j|i}$ (for ground truth) and $d_{j|i}$ (for constrained p-correlation), and quantities analogous to $d_{j|i}$ (for Granger B1, Gen Synch S1, LiNGAM, and Patel) for Subject 14 of $Sim2$. Each image uses the same ordering of colors, but has different range of numerical values.

The mean and standard deviation of accuracy for each simulation, i.e., the average and square root of the sample variance of $\mathcal{A}$ (Eq. 6) over all 50 subjects, were computed and the results are tabulated in Table 1. For all four simulations, constrained p-correlation achieved the highest accuracy compared to other methods. The unconstrained p-correlation is less appropriate when applied to a network with all positive connection weights. We also computed the mean and standard deviation of $\mathcal{A}$ for pairwise LiNGAM, which gives .566±.138, .656±.206, .510±.119 and .506±.056 for four simulations, respectively. The result shows the highly accurate directionality that pairwise LiNGAM can achieve in this particular unidirectional network setting. Histograms displaying the distribution of accuracy for the five methods for each simulation are shown in Fig. 5. The superior performance of p-correlation is demonstrated by the fact that the bulk of the histogram is further to the right, and the left tail is less massive.

Simulation	1	2	3	4
# of ROIs	5	10	15	50
# of COI pairs	10	22	36	122
Granger B1	.440±.206	.295±.127	.262±.088	.130±.044
Gen Synch S1	.472±.201	.405±.139	.379±.079	.285±.056
LiNGAM	.372±.229	.435±.177	.301±.106	.119±.037
Patel	.528±.193	.491±.101	.446±.099	.366±.048
p-Corr (constrained)	.532±.192	.502±.114	.457±.126	.405±.065
p-Corr (unconstrained)	.520±.218	.467±.123	.439±.109	.371±.058

Table 1: Comparison of the mean and standard deviation of accuracy over 50 subjects among different methods. The standard deviation tend to be larger on the smaller netowrks ($Sim1$-$Sim4$ have 5, 10, 15, 50 ROIs, respectively) because one error is proportionally of larger impact in a smaller network.

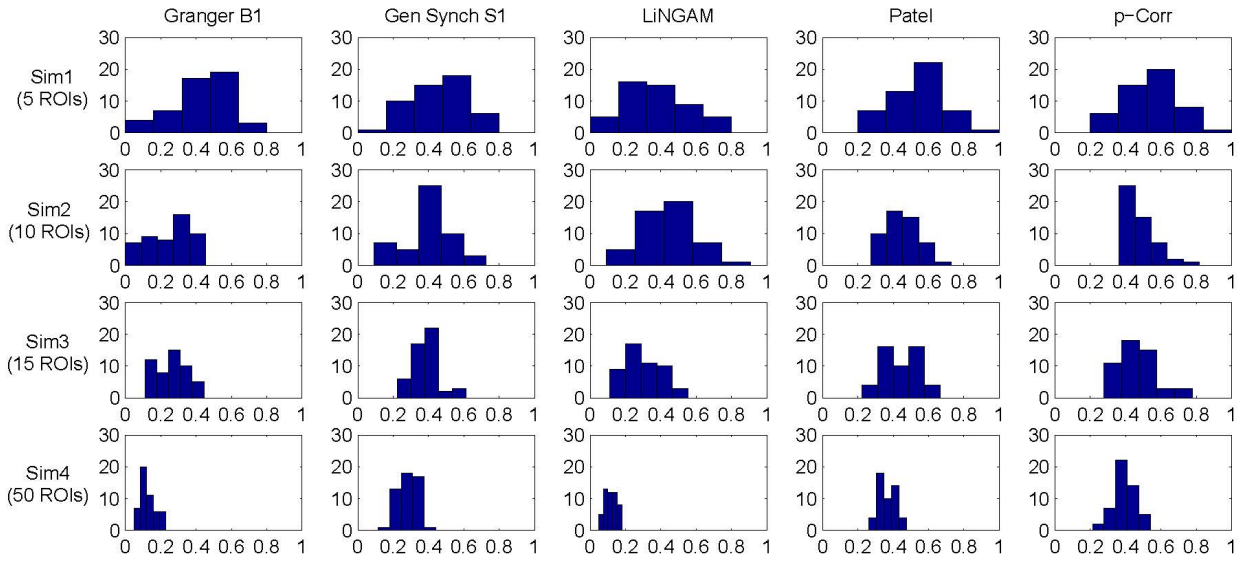

Figure 5: Accuracy histogram for Granger B1, Gen Synch S1, LiNGAM, Patel and constrained p-correlation.

3.2. The performance of correlation and p-correlation on common drivers

A "common driver" situation is the case where ROI 1 drives ROIs 2 and 3 but ROIs 2 and 3 do not directly interact. The challenge is to correctly detect the 1→2 and 1→3 connections without detecting 2→3 or 3→2 false connections. In order to focus exclusively on this situation, we have computed synthetic data from the three-ROI network shown in Fig. 6 and defined by

$$x_1[n+1] = a_1 x_1[n] + b_1 w_1[n] \tag{7}$$

$$x_2[n+1] = a_2 x_2[n] + a_{21} x_1[n] + b_2 w_2[n] \tag{8}$$

$$x_3[n+1] = a_3 x_3[n] + a_{31} x_1[n] + b_3 w_3[n] \tag{9}$$

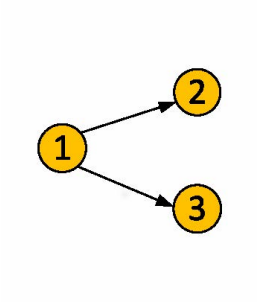

Figure 6: The common driver problem.

where $w[n] = [w_1[n], w_2[n], w_3[n]]^T$ is an independent and identically distributed Gaussian stochastic process with mean 0 and variance I_3 (the 3×3 identity matrix). Zalesky et al. (2012) consider mathematical models of this type and give theoretical results for correlations. The system is initialized in the steady state and simulated for 1000 steps, $N_x = 1000$. We consider only $a_1 = a_2 = a_3 = .8$ (so that all ROIs have the same intrinsic memory duration) and $b_1 = b_2 = b_3 = .2$ (so that all ROIs have the same intrinsic noise power, and the intrinsic noises are all independent). We consider the following cases: (1) no driving: $a_{21} = a_{31} = 0$, (2) weak driving: $a_{21} = a_{31} = .1$, (3) strong driving: $a_{21} = a_{31} = .4$, and (4) asymmetrical strong driving: $a_{21} = .4$, and $a_{31} = .1$.

Each simulation was repeated for 50 subjects. Let the maximum allowable duration of the impulse response be 3 samples. By using the specialization of p-correlation for Smith simulated data, as is described in Section 3.1.2, a directed graph $d_{j|i}$ is estimated by p-correlation (Fig. 2) and the correlation matrix is computed for each subject. The steady

state covariance of Eqs. 7-9 is the correlation matrix. In Case (1), the mean and standard deviation of nonzero entries of $\rho_{j|i}$ with constrained least squares (Section 2.2.1) are 5.384e-04$\pm$0.072. This number becomes 0.058$\pm$0.043 when unconstrained least squares is applied. The smaller magnitude of the results using constrained least squares indicates that taking advantage of the prior knowledge that the weights are positive (i.e., $a_1 = a_2 = a_3 = .8$) provides improved performance in this case. In Cases (2) and (3), both the constrained and the unconstrained least squares achieve a 100% accuracy (Eq. 6) for each subject. In the fourth case, the constrained or the unconstrained least squares gives an average of .800$\pm$.247 accuracy over all 50 subjects. We also tested N_x =200, 500, 5000 for all four cases. Notice that as N_x goes large, correlations become closer to the steady state and the accuracy computed by the p-correlation method increases as well.

In addition, p-correlation estimated the correct hierarchy on the three pairs of connection weights, which are consistent with "strong", "weak" and "non-" connections in the ground truth network. It also shows the correct direction of connections in a pair by a stronger weight. The constrained least squares (Section 2.2.1) provides a slightly superior result than the unconstrained approach. Specifically, larger numerical differences between the zero and nonzero entries, as well as between the asymmetric strong weights, were shown. On average across all 50 subjects, p-correlation used an impulse response duration of 1.007 samples for all four cases for both constrained and unconstrained approaches. In addition, in Case (3) (asymmetric strong weights), correlation mis-detected the connection between node 2 and 3, specifically the 2-3 correlation was the highest correlation value among the three pairs, whereas p-correlation, for both the constrained and unconstrained approaches, estimated this value as the lowest of the three pairs thereby avoiding the error in the correlation results.

3.3. Performance on experimental fMRI data

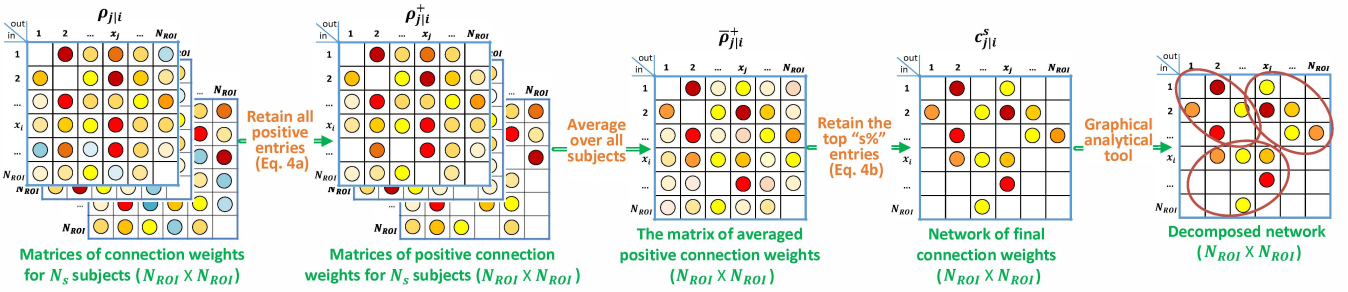
Retain all positive entries (Eq. 4a)
Average over all subjects
Retain the top "s%" entries (Eq. 4b)
Graphical analytical tool
Matrices of connection weights for N_s subjects (N_{ROI} X N_{ROI})
Matrices of positive connection weights for N_s subjects (N_{ROI} X N_{ROI})
The matrix of averaged positive connection weights (N_{ROI} X N_{ROI})
Network of final connection weights (N_{ROI} X N_{ROI})
Decomposed network (N_{ROI} X N_{ROI})

Figure 7: Block diagram describing the specialization of p-correlation for the experimental data. Nonzero entries are filled by colored dots with higher values represented by "hotter" colors and lower values represented by "colder" colors, and zero entries are left as blank in the above matrices.

While the tools described in this paper can be assembled into many algorithms, we use only one algorithm, which is shown in Figure 7, to further characterize (Xu et al., 2014), a cohort of 132 subjects from the 1000 Functional Connectomes Project (`http://www.nitrc.org/projects/fcon_1000/`) (Biswal et al., 2010). This data is provided from different scanning sites, and thus has variable sampling rates (TRs = approx 1-3s, mean ± standard deviation of 2.3 ± 0.4s). The scan duration also varied from 119-295 TRs (mean standard deviation of 167.5 ± 41.7). The data from the whole brain were preprocessed (Anderson et al., 2011), linearly detrended and bandpass filtered (retaining signal between 0.001 and 0.1 Hz), and motion scrubbed (Power et al., 2012) with the threshold set to 0.2. The preprocessed rs-fMRI BOLD signal was extracted from $N_{\mathrm{ROI}} = 264$ spherical ROIs each with a 10mm diameter. We combine our p-correlation ideas with the widely-used (Power et al., 2011, 2012; Lahnakoski et al., 2012; Gordon et al., 2016) Infomap graph analytical algorithm (Lancichinetti and Fortunato, 2009) to determine networks within the set of 264 ROIs.

As a function of the value of the threshold s, Infomap creates a variable number of networks. Following Power et al. (2011, Fig. 1), the network stability over a range of threshold $s \in \{2, \ldots, 10\}$ using correlations and p-correlations are shown in Fig. 8, in which different networks are represented by different colors. Similar to Power et al. (2011) (the first figure in Fig 8), we note that the assignment of ROIs to networks remains relatively

constant over all values of the threshold s, illustrated by the constant horizontal bands in different colors. Also, networks are hierarchically refined as s rises. In summary, the number of networks increases as the value of s decreases, and p-correlation replicated the brain network organizations that were detected by correlation. The network results are consistent with the network organizations detected in Power et al. (2011).

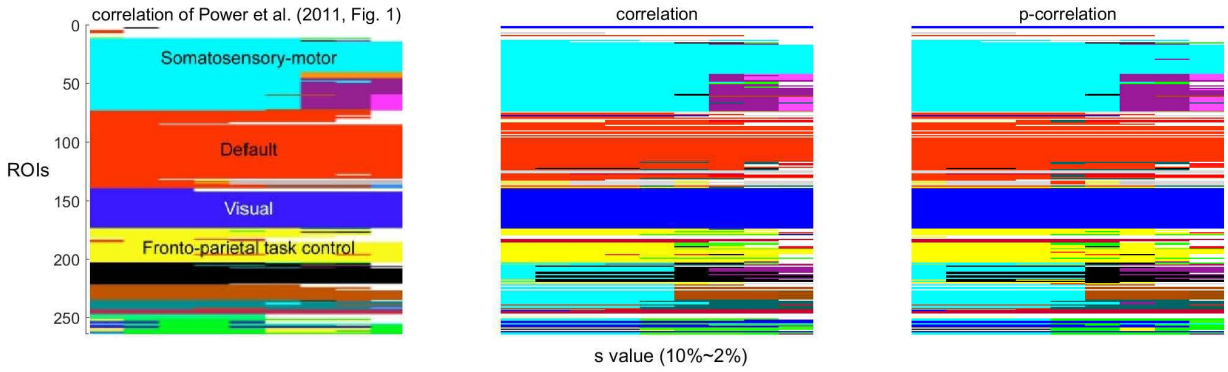

Figure 8: The stability of networks across various thresholding criteria (s). The white regions indicate ROIs that belong to networks with less than four ROIs.

In order to test the robustness of the p-correlation calculation, all 132 subjects were randomly divided into two equal cohorts, and each cohort was separately processed. The average of p-correlation connection strength $\rho^+_{j|i}$ across all subjects in the cohort, which is denoted by $\bar{\rho}^+_{j|i}$, is shown as a scatter plot in Fig. 9 (a) (in Fig. 9, all (0,0) points are removed). The linear least squares prediction of Cohort 2 from Cohort 1 is a close fit to the data ($r^2 = .87$) and is nearly a 45° diagonal line ($\bar{\rho}^{\text{Cohort 2}}_{j|i}$=1.013$\bar{\rho}^{\text{Cohort 1}}_{j|i}$+.032), thereby indicating the robust nature of p-correlation. Following the same procedure, the average of correlation connection strength $\rho^+_{i,j}$ across all subjects in the cohort, which is denoted by $\bar{\rho}^+_{i,j}$, is shown in Fig. 9 (b). Comparing Fig.9 (a) and (b) indicates that the p-correlation achieves the same robustness as correlation. Additional plots in which no points are removed are included in the supplemental material Fig. 2.

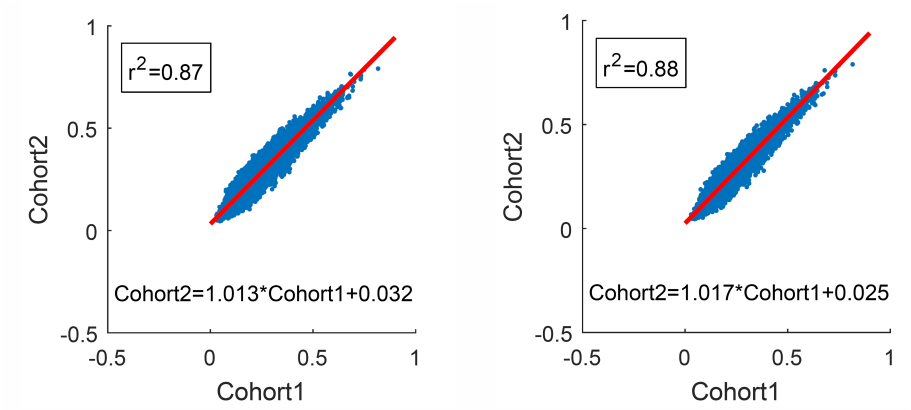

(a) p-correlation (b) correlation

Figure 9: Scatter plot of p-correlation and correlation for the two cohorts. The red line is the Least Squares fit for predicting Cohort 2 from Cohort 1. Only positive values are used in the Least Squares calculation and shown in the plot.

4. Discussion

Standard correlation has been widely used to analyze functional connectivity from rs-fMRI timeseries between prespecified ROIs. Prior work has shown its high sensitivity for detecting the existence of network architectures under both simulated and experimental scenarios (Smith et al., 2011; Dawson et al., 2013). This paper describes methodology for analyzing rs-fMRI data using a generalization of well-established correlation ideas. The generalization, denoted by "p-correlation" ("p" for "prediction"), is to compute the correlation between the j^{th} signal and an optimal linear time-invariant causal estimate of the j^{th} signal based on the i^{th} signal. In this way, it captures additional features concerning the interaction between two ROIs, specifically, the causality and directionality of the information flow on which the interaction depends. Based on the finite-memory linear time-invariant causal model, p-correlation allows the memory duration to be different in the two directions for one pair of ROIs and also to be different for different pairs of ROIs. In contrast, structural vector autoregressive models (Kim et al., 2007; Chen et al., 2011) are assumed to have the same memory duration across all ROIs. P-correlation is a generalization of standard correlation ideas because, if the estimate of the j^{th} signal based on the i^{th} signal is restricted

to use only the current value of the i^{th} signal, then p-correlation and standard correlation have the same magnitude.

Testing p-correlation on simulated fMRI data provided in Smith et al. (2011), the greater performance accuracy of p-correlation, which uses lagged information from the BOLD time-series, demonstrates the importance of causal information which is missing in standard correlation. In our results, the mean duration of the impulse response estimated by AIC using a search limited to a maximum duration of 15s was roughly 4s. In these data, a search extending to 15s is not a restriction on the maximum duration. As is described in Table 1, the accuracy of p-correlation on the simulated data of Smith is about 0.5 (0.405 to 0.532). While higher levels are desirable, this performance exceeds the performance of many alternative algorithms on all four sets of simulations.

Many approaches have been introduced to assess functional or effective connectivity of rs-fMRI data. Smith et al. (2011) evaluated the validity of 38 approaches (Smith et al., 2011, Fig. 4) using simulated BOLD signals and a variety of performance measures. The methods tend to have different levels of performance for different measures, e.g., detection of a connection versus determination of the direction of a connection. The p-correlation approach introduced in this paper depends on causal dynamical models and so we focus on this particular aspect of previous work. Dynamic Causal Modeling (DCM) has been used with some success to assess causal dynamics in fMRI data by relying on sophisticated models of neural dynamics. As discussed in Smith et al. (2011, p. 878), most existing DCM algorithms require knowledge of external inputs (which are not known for rs-fMRI) although some variations may not (Daunizeau et al., 2009); all versions tend to be mathematically poorly conditioned; and all versions fail to scale to networks with large numbers of ROIs which are necessary for experimental studies. In contrast, the p-correlation approach described in this paper scales similarly to a correlation approach for which hundreds of ROIs are not a challenge (Xu et al., 2014).

Several versions of Granger causality analysis, based on multivariate vector autoregressive modeling, have been tested and performed poorly (Smith et al., 2011). Granger causality relies on regression and comparison of two predictions. The first prediction is based purely on an autoregressive model of the signal at the ith ROI based on the past of the same signal. The second prediction is based on regression of the signal at the ith ROI based on the past of the signal at the jth ROI and, possibly, an autoregression as in the first case. The sample

covariances of the prediction errors are then combined, essentially by taking the ratio of the sample covariances scaled by integers describing the amounts of data, to yield a statistic that is distributed according to the Fisher-Snedecor F distribution. This statistic, indexed by i and j, is used to fill an asymmetric matrix. Although both are based upon lagged information there are important differences between p-correlation and Granger causality. P-correlation is not a statistic comparing two possible dependencies but rather is a statistic measuring the accuracy of prediction using a particular dependency. The motivation for the Granger causality statistic is dependent on the original Gaussian assumptions on the errors when linear regression is used to describe the ROI time series. P-correlation is based on just the sample variance of the prediction error and does not have a Gaussian motivation which is advantageous if the BOLD signals lack Gaussian structure. Multivariate autoregressive processes have been used as the basis for generative models for complete sets of ROIs. Such models, which focus on the effect of the past on the present, can be combined with structural equation modeling (SEM) models, which focus on contemporaneous effects (Chen et al., 2011).

Multivariate autoregressive processes (MVAR) have been successfully used in neuroscience outside of fMRI, e.g., in order to describe signals from EEG experiments (Kus et al., 2004; Blinowska et al., 2009; Ding et al., 2000; Babiloni et al., 2005; Ligeza et al., 2016; Korzeniewska et al., 2011; Wilke et al., 2008). Both MVAR, e.g., Eq. 1 in Kus et al. (2004), and the linear regression model used in this paper (Eq. 1) are regression models which predict one timeseries from either all timeseries which include oneself (MVAR) or from the past of another timeseries (Eq. 1). Both predictions are characterized by impulse responses. The method introduced in Kus et al. (2004) determines the connection strength based on the impulse response, whereas p-correlation determines the functional connectivity based on both the impulse response and the original timeseries. Existing literature, e.g., Valdés-Sosa et al. (2005) and Davis et al. (2015), has shown the robust estimation of the MVAR model by introducing sparse regression techniques, and the success of estimating functional connectivity through the sparse MVAR models. In addition, a conditional MVAR model, e.g., Ch 17.3 in Schelter et al. (2006), may also be used to address the common driver problem. Other approaches to examining BOLD signal propagation using lags, as is done in p-correlation, have been highly reproducible (Mitra et al., 2015). In this paper, a linear regression model (Eq. 1) is used as the predictor in p-correlation to estimate the causal

relation between a pair of BOLD signals. Other lag-based predictors, e.g., MVAR based models, can also be adapted into the p-correlation concept, however, they would not have the result that duration of 1 sample (e.g., no lags) gives standard correlation.

In addition to the algorithms used in Smith et al. (2011), which estimate the directional connectivity for single subject data sets, the IMaGES (Ramsey et al., 2010, 2011) and GIMME (Gates and Molenaar, 2012) algorithms use a group of subjects. While these algorithms provide better performance in situations where groups of subjects can be analyzed collectively, both algorithms have challenges. The sparse graph estimated by IMaGES for a group of subjects does not tell the strengths of the connectivity and "will not reflect the variation of a group" (Mumford and Ramsey, 2014, p.571). Similar to DCM's limitation on scalability, small networks with less than 25 ROIs are well analyzed by the GIMME algorithm. However, its performance on large-scale functional networks is not known. As p-correlation can work with hundreds of ROIs, it can be used in evaluating large-scale brain networks. Furthermore, p-correlation can work on individual subjects so it potentially could be applied to patient clinical data. Other algorithms that estimate direction after a connection is already detected also exist (Section 3.1.4). While such algorithms may be useful in some circumstances, they do not allow for situations where both directions are present but of different strengths.

The Smith et al. (2011) simulated data has lower dimensionality than experimental brain data. For instance, in the simulation, connections are all unidirectional while most neural connections are bidirectional. Additionally, in the simulations, most connections had a value of exactly zero. Furthermore, it introduces unrealistic noise and it has a large number of parameters that must be set and which influence the resulting simulation (Wang et al., 2014). While the Smith et al. (2011) simulated data is not completely realistic and it is not a perfect test of p-correlation, this data continues to be used (Smith et al., 2011; Gates and Molenaar, 2012; Ramsey et al., 2011; Hyvärinen and Smith, 2013; Ryali et al., 2016), and the results continue to be discussed (Geerligs et al., 2016). In this paper, we leveraged the same data used in Smith et al. (2011) for comparison with other published metrics, providing a broader context for these findings. We hope to use a broader range of simulated data to further validate p-correlation in our future work.

In order to focus on the challenges of a "common driver", we have produced additional synthetic data for the three ROI network of Fig. 6 in which one ROI drives two other ROIs

but the two other ROIs do not directly interact. Using p-correlation in this network we found that p-correlation can identify the existence and direction of the interactions between the driving ROI and the other two ROIs (even when the two interactions are of different strengths). Furthermore, p-correlation did not introduce false interactions between the two driven ROIs.

We have applied p-correlation to experimental data from the 1000 Functional Connectome Project (Biswal et al., 2010). The p-correlation approach successfully replicated the modular architecture of the local and distributed networks previously reported using standard correlation (Xu et al., 2014) (see Section 3.3 Fig. 8). Highly correlated p-correlation values on the two different cohorts also demonstrated that the p-correlation is highly reproducible and thus robust on experimental data. A current limitation of the p-correlation approach is that missing nodes cannot be accommodated, thereby limiting an extension of this approach to lesioned populations.

Here we introduce a novel concept, the p-correlation, to estimate brain connectivity within well-characterized large-scale functional networks. The replication of previously observed network architectures in experimental data and the performance against the ground truth in simulated data, both suggest that the p-correlation approach may hold promise for future investigations of the brain's dynamic functional architecture.

5. Acknowledgements

We are very grateful to Prof. S.M. Smith (University of Oxford) for providing simulation data and his software for applying Patel's conditional dependence measures and network measurements as described in his paper (Smith et al., 2011) as well as for helpful discussion. We are also very grateful to Dr. Chandler Lutz for providing the pairwise Granger causality code, Dr. Rodrigo Quian Quiroga for providing the generalized synchronization code, and Drs. Shohei Shimizu, Patrik Hoyer and Aapo Hyvrinen for providing LiNGAM/FastICA. Nan Xu and Peter C. Doerschuk are grateful for support by NSF Grant 1217867.

Supplemental

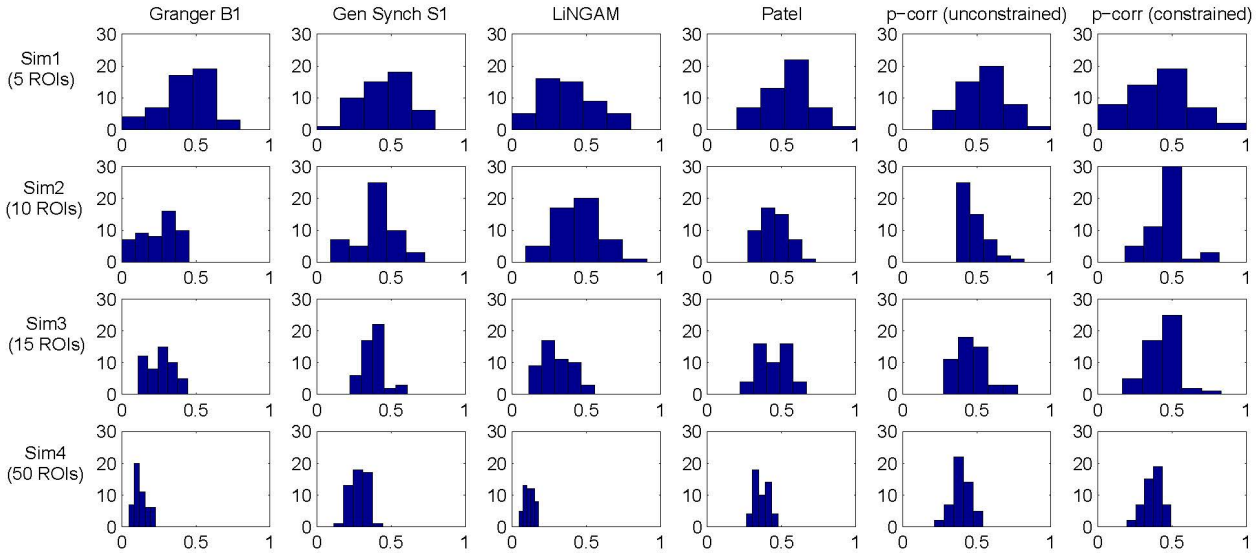

Figure S1: Accuracy histogram for Granger B1, Gen Synch S1, LiNGAM, Patel and p-correlation with unconstrained and constrained Least Squares.

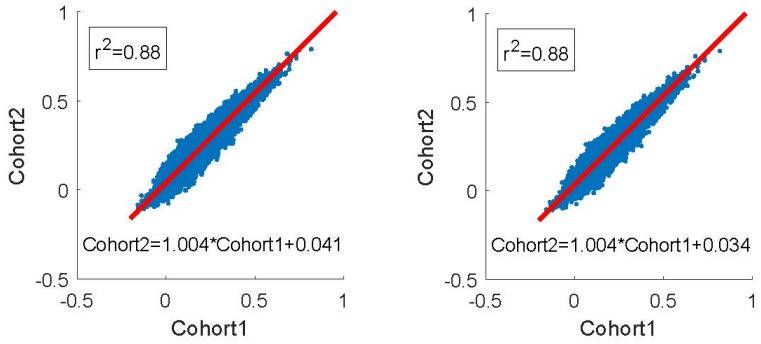

(a) p-correlation (b) correlation

Figure S2: Scatter plot of p-correlation and correlation for the two cohorts. The red line is the Least Squares fit for predicting Cohort 2 from Cohort 1.

References

Akaike, H., 1970. Statistical predictor identification. Annals of the Institute of Statistical Mathematics 22, 203–217.

Akaike, H., 1974. A new look at the statistical model identification. IEEE Transactions on Automatic Control 19, 716–723.

Anderson, J.S., Druzgal, T.J., Lopez-Larson, M., Jeong, E.K., Desai, K., Yurgelun-Todd, D., 2011. Network anticorrelations, global regression, and phase-shifted soft tissue correction. Human brain mapping 32, 919–934.

Babiloni, F., Cincotti, F., Babiloni, C., Carducci, F., Mattia, D., Astolfi, L., Basilisco, A., Rossini, P., Ding, L., Ni, Y., et al., 2005. Estimation of the cortical functional connectivity with the multimodal integration of high-resolution eeg and fmri data by directed transfer function. Neuroimage 24, 118–131.

Bassett, D.S., Gazzaniga, M.S., 2011. Understanding complexity in the human brain. Trends in cognitive sciences 15, 200–209.

Biswal, B., Zerrin Yetkin, F., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magnetic resonance in medicine 34, 537–541.

Biswal, B.B., Mennes, M., Zuo, X.N., Gohel, S., Kelly, C., Smith, S.M., Beckmann, C.F., Adelstein, J.S., Buckner, R.L., Colcombe, S., et al., 2010. Toward discovery science of human brain function. Proceedings of the National Academy of Sciences 107, 4734–4739.

Blinowska, K., Trzaskowski, B., Kaminski, M., Kus, R., 2009. Multivariate autoregressive model for a study of phylogenetic diversity. Gene 435, 104–118.

Cavanaugh, J.E., 1997. Unifying the derivations for the Akaike and corrected Akaike information criteria. Statistics & Probability Letters 33, 201–208.

Chen, G., Glen, D.R., Saad, Z.S., Hamilton, J.P., Thomason, M.E., Gotlib, I.H., Cox, R.W., 2011. Vector autoregression, structural equation modeling, and their synthesis in neuroimaging data analysis. Computers in biology and medicine 41, 1142–1155.

Cole, M.W., Bassett, D.S., Power, J.D., Braver, T.S., Petersen, S.E., 2014. Intrinsic and task-evoked network architectures of the human brain. Neuron 83, 238–251.

Daunizeau, J., Friston, K., Kiebel, S., 2009. Variational bayesian identification and prediction of stochastic nonlinear dynamic causal models. Physica D: Nonlinear Phenomena 238, 2089 – 2118. doi:10.1016/j.physd.2009.08.002.

Dauwels, J., Vialatte, F., Musha, T., Cichocki, A., 2010. A comparative study of synchrony measures for the early diagnosis of Alzheimer's disease based on eeg. NeuroImage 49, 668–693.

Davis, R.A., Zang, P., Zheng, T., 2015. Sparse vector autoregressive modeling. Journal of Computational and Graphical Statistics , 1–53.

Dawson, D.A., Cha, K., Lewis, L.B., Mendola, J.D., Shmuel, A., 2013. Evaluation and calibration of functional network modeling methods based on known anatomical connections. NeuroImage 67, 331–343.

Ding, M., Bressler, S.L., Yang, W., Liang, H., 2000. Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment. Biological cybernetics 83, 35–45.

Foster, B.L., Rangarajan, V., Shirer, W.R., Parvizi, J., 2015. Intrinsic and task-dependent coupling of neuronal population activity in human parietal cortex. Neuron 86, 578–590. doi:10.1016/j.neuron.2015.03.018.

Friston, K.J., 2011. Functional and effective connectivity: A review. Brain Connectivity 1, 13–36.

Gates, K.M., Molenaar, P.C., 2012. Group search algorithm recovers effective connectivity maps for individuals in homogeneous and heterogeneous samples. Neuroimage 63, 310–319.

Geerligs, L., Henson, R.N., et al., 2016. Functional connectivity and structural covariance between regions of interest can be measured more accurately using multivariate distance correlation. NeuroImage 135, 16–31.

Gordon, E.M., Laumann, T.O., Adeyemo, B., Huckins, J.F., Kelley, W.M., Petersen, S.E., 2016. Generation and evaluation of a cortical area parcellation from resting-state correlations. Cerebral cortex 26, 288–303.

Hipp, J.F., Siegel, M., 2015. Bold fMRI correlation reflects frequency-specific neuronal correlation. Current Biology 25, 1368–1374.

Hurvich, C.M., Tsai, C.L., 1989. Regression and time series model selection in small samples. Biometrika 76, 297–307.

Hurvich, C.M., Tsai, C.L., 1993. A corrected Akaike information criterion for vector autoregressive model selection. Journal of time series analysis 14, 271–279.

Hyvärinen, A., Smith, S.M., 2013. Pairwise likelihood ratios for estimation of non-gaussian structural equation models. The Journal of Machine Learning Research 14, 111–152.

Jo, H.J., Gotts, S.J., Reynolds, R.C., Bandettini, P.A., Martin, A., Cox, R.W., Saad, Z.S., 2013. Effective preprocessing procedures virtually eliminate distance-dependent motion artifacts in resting state fmri. J. Applied Mathematics 2013, 935154:1–935154:9.

Kim, J., Zhu, W., Chang, L., Bentler, P.M., Ernst, T., 2007. Unified structural equation modeling approach for the analysis of multisubject, multivariate functional MRI data. Human Brain Mapping 28, 85–93.

Korzeniewska, A., Franaszczuk, P.J., Crainiceanu, C.M., Kuś, R., Crone, N.E., 2011. Dynamics of large-scale cortical interactions at high gamma frequencies during word production: event related causality (erc) analysis of human electrocorticography (ecog). Neuroimage 56, 2218–2237.

Kus, R., Kaminski, M., Blinowska, K.J., 2004. Determination of eeg activity propagation: pair-wise versus multichannel estimate. IEEE transactions on Biomedical Engineering 51, 1501–1510.

Lahnakoski, J.M., Glerean, E., Salmi, J., Jääskeläinen, I.P., Sams, M., Hari, R., Nummenmaa, L., 2012. Naturalistic fMRI mapping reveals superior temporal sulcus as the hub for the distributed brain network for social perception. Frontiers in human neuroscience 6.

Lancichinetti, A., Fortunato, S., 2009. Community detection algorithms: A comparative analysis. Phys. Rev. E 80, 056117. doi:10.1103/PhysRevE.80.056117.

Laumann, T.O., Gordon, E.M., Adeyemo, B., Snyder, A.Z., Joo, S.J., Chen, M.Y., Gilmore, A.W., McDermott, K.B., Nelson, S.M., Dosenbach, N.U., et al., 2015. Functional system and areal organization of a highly sampled individual human brain. Neuron 87, 657–670.

Ligeza, T.S., Wyczesany, M., Tymorek, A.D., Kamiński, M., 2016. Interactions between the prefrontal cortex and attentional systems during volitional affective regulation: an effective connectivity reappraisal study. Brain topography 29, 253–261.

Mitra, A., Snyder, A.Z., Blazey, T., Raichle, M.E., 2015. Lag threads organize the brain's intrinsic activity. Proceedings of the National Academy of Sciences of the United States of America 112, E2235–E2244. doi:10.1073/pnas.1503960112.

Mumford, J.A., Ramsey, J.D., 2014. Bayesian networks for fMRI: a primer. Neuroimage 86, 573–582.

Murphy, K., Birn, R.M., Handwerker, D.A., Jones, T.B., Bandettini, P.A., 2009. The impact of global signal regression on resting state correlations: Are anti-correlated networks introduced? NeuroImage 44, 893–905.

Patterson, H.D., Thompson, R., 1971. Recovery of inter-block information when block sizes are unequal. Biometrika 58, 545–554.

Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage 59, 2142–2154.

Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., et al., 2011. Functional network organization of the human brain. Neuron 72, 665–678.

Power, J.D., Schlaggar, B.L., Lessov-Schlaggar, C.N., Petersen, S.E., 2013. Evidence for hubs in human functional brain networks. Neuron 79, 798–813.

Ramsey, J.D., Hanson, S.J., Glymour, C., 2011. Multi-subject search correctly identifies causal connections and most causal directions in the dcm models of the smith et al. simulation study. NeuroImage 58, 838–848.

Ramsey, J.D., Hanson, S.J., Hanson, C., Halchenko, Y.O., Poldrack, R.A., Glymour, C., 2010. Six problems for causal inference from fMRI. Neuroimage 49, 1545–1558.

Rissanen, J., 1978. Modeling by shortest data description. Automatica 14, 465–471.

Ryali, S., Chen, T., Supekar, K., Tu, T., Kochalka, J., Cai, W., Menon, V., 2016. Multivariate dynamical systems-based estimation of causal brain interactions in fmri: Group-level validation using benchmark data, neurophysiological models and human connectome project data. Journal of neuroscience methods 268, 142–153.

Saad, Z.S., Gotts, S.J., Murphy, K., 0008, G.C., Jo, H.J., Martin, A., Cox, R.W., 2012. Trouble at rest: How correlation patterns and group differences become distorted after global signal regression. Brain Connectivity 2, 25–32.

Sadaghiani, S., Poline, J.B., Kleinschmidt, A., DEsposito, M., 2015. Ongoing dynamics in large-scale functional connectivity predict perception. Proceedings of the National Academy of Sciences 112, 8463–8468.

Schaefer, A., Margulies, D.S., Lohmann, G., Gorgolewski, K.J., Smallwood, J., Kiebel, S.J., Villringer, A., 2014. Dynamic network participation of functional connectivity hubs assessed by resting-state fMRI. Frontiers in human neuroscience 8.

Schelter, B., Winterhalder, M., Timmer, J., 2006. Handbook of time series analysis: recent theoretical developments and applications. John Wiley & Sons.

Schwarz, G., 1978. Estimating the dimension of a model. The Annals of Statistics 6, 461–464.

Smith, S.M., 2012. The future of fmri connectivity. Neuroimage 62, 1257–1266.

Smith, S.M., Miller, K.L., Salimi-Khorshidi, G., Webster, M., Beckmann, C.F., Nichols, T.E., Ramsey, J.D., Woolrich, M.W., 2011. Network modelling methods for fMRI. NeuroImage 54, 875–891.

Sporns, O., 2011. The human connectome: a complex network. Annals of the New York Academy of Sciences 1224, 109–125.

Sugiura, N., 1978. Further analysis of the data by Akaike's information criterion and the finite corrections. Communications in Statistics - Theory and Methods 7, 13–26. doi:10.1080/03610927808827599.

Thompson Jr, W., 1962. The problem of negative estimates of variance components. The Annals of Mathematical Statistics , 273–289.

Valdés-Sosa, P.A., Sánchez-Bornot, J.M., Lage-Castellanos, A., Vega-Hernández, M., Bosch-Bayard, J., Melie-García, L., Canales-Rodríguez, E., 2005. Estimating brain functional connectivity with sparse multivariate autoregression. Philosophical Transactions of the Royal Society B: Biological Sciences 360, 969–981.

Van Den Heuvel, M.P., Pol, H.E.H., 2010. Exploring the brain network: a review on resting-state fMRI functional connectivity. European Neuropsychopharmacology 20, 519–534.

Wallace, C.S., Boulton, D.M., 1968. An information measure for classification. The Computer Journal 11, 185–194.

Wang, H.E., Bénar, C.G., Quilichini, P.P., Friston, K.J., Jirsa, V.K., Bernard, C., 2014. A systematic framework for functional connectivity measures. Frontiers in neuroscience 8, 405.

Wig, G., Schlaggar, B., Petersen, S., 2011. Concepts and principles in the analysis of brain networks. Annals of the New York Academy of Sciences 1224, 126–146. doi:10.1111/j.1749-6632.2010.05947.x.

Wilke, C., Ding, L., He, B., et al., 2008. Estimation of time-varying connectivity patterns through the use of an adaptive directed transfer function. IEEE transactions on biomedical engineering 55, 2557–2564.

Xu, N., Spreng, R.N., Doerschuk, P.C., 2014. Directed interactivity of large-scale brain networks: introducing a new method for estimating resting-state effective connectivity MRI, in: Image Processing (ICIP), 2014 21st IEEE International Conference on, pp. 3508–3512.

Yeo, B.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., et al., 2011. The organization

of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of neurophysiology 106, 1125–1165.

Zalesky, A., Fornito, A., Bullmore, E., 2012. On the use of correlation as a measure of network connectivity. Neuroimage 60, 2096–2106.